

Hello and thanks for reading. This is fairly random collection of text, mostly from Wikipedia, that I came across as I was studying several topics related to CISSP. I did copy/pastes here, with very little editing. Much of the relevant data from this document, regarding CISSP certification, is already included in the CISSP course from CBTNuggets.com, so enjoying those videos would be a great way to prepare for the CISSP certification. For any hyperlinked fields that may appear in this document, I didn't verify each one, and they may change over time. As a result, I would encourage you to go directly to the source (using Google, Wikipedia, or other Internet resource directly from your browser) for additional detail, instead of using any links which may be in this document. Best wishes in your CISSP studies.

Enjoy!

Risk management is the identification, assessment, and prioritization of [risks](#) (defined in [ISO 31000](#) as *the effect of uncertainty on objectives*) followed by coordinated and economical application of resources to minimize, monitor, and control the probability and/or impact of unfortunate events^[1] or to maximize the realization of opportunities. Risk management's objective is to assure [uncertainty](#) does not deflect the endeavor from the business goals.^[2]

[Electronic discovery](#) or "e-discovery" refers to discovery of information stored in [electronic format](#) (often referred to as [Electronically Stored Information](#), or ESI).

A **Security Token Service (STS)** is a software based [identity provider](#) responsible for issuing [security tokens](#), especially [software tokens](#), as part of a [claims-based identity](#) system.^[1]

In a typical usage scenario, a client requests access to a secure software application, often called a relying party. Instead of the application authenticating the client, the client is redirected to an STS. The STS authenticates the client and issues a security token. Finally, the client is redirected back to the relying party where it presents the security token. The token is the data record in which claims are packed, and is protected from manipulation with strong cryptography. The software application verifies that the token originated from an STS trusted by it, and then makes authorization decisions accordingly. The token is creating a chain of trust between the STS and the software application consuming the claims. This process is illustrated in the [Security Assertion Markup Language \(SAML\)](#) use case, demonstrating how single sign-on can be used to access web services.

OAuth is an [open standard](#) for [authorization](#). OAuth provides client applications a 'secure delegated access' to server resources on behalf of a resource owner. It specifies a process for resource owners to authorize third-party access to their server resources without sharing their credentials. Designed specifically to work with [Hypertext Transfer Protocol](#) (HTTP), OAuth essentially allows [access tokens](#) to be issued to third-party clients by an authorization server, with the approval of the resource owner. The client then uses the access token to access the protected resources hosted by the resource server.^[1] OAuth is commonly used as a way for Internet users to log into third

party websites using their Microsoft, Google, Facebook or Twitter accounts without exposing their password.^[2]

OpenID is an [open standard](#) and [decentralized protocol](#) by the non-profit **OpenID Foundation** that allows users to be [authenticated](#) by certain co-operating sites (known as [Relying Parties or RP](#)) using a third party service. This eliminates the need for [webmasters](#) to provide their own ad hoc systems and allowing users to consolidate their [digital identities](#).^[1]

Disaster recovery (DR) involves a set of policies and procedures to enable the recovery or continuation of vital technology infrastructure and systems following a [natural](#) or [human-induced disaster](#).^[1] Disaster recovery focuses on the IT or [technology systems](#) supporting critical business functions,^[2] as opposed to [business continuity](#), which involves keeping all essential aspects of a business functioning despite significant disruptive events. Disaster recovery is therefore a subset of [business continuity](#).^[3]

A **penetration test**, or sometimes **pentest**, is a software attack on a computer system that looks for security weaknesses, potentially gaining access to the computer's features and data.^{[1][2]}

The process typically identifies the target systems and a particular goal—then reviews available information and undertakes various means to attain the goal. A penetration test target may be a [white box](#) (which provides background and system information) or [black box](#) (which provides only basic or no information except the company name). A penetration test can help determine whether a system is vulnerable to attack, if the defenses were sufficient, and which defenses (if any) the test defeated.^[3]

Trusted Platform Module (TPM) is an [international standard](#) for a [secure cryptoprocessor](#), which is a dedicated microprocessor designed to secure hardware by integrating cryptographic keys into devices.

Social engineering, in the context of [information security](#), refers to [psychological manipulation](#) of people into performing actions or divulging confidential information. A type of [confidence trick](#) for the purpose of information gathering, fraud, or system access, it differs from a traditional "con" in that it is often one of many steps in a more complex fraud scheme.

The term "social engineering" as an act of psychological manipulation is also associated with the social sciences, but its usage has caught on among computer and information security professionals.^[1]

Certification and Accreditation (C&A or CnA) is a process for implementing any formal process. It is a systematic procedure for evaluating, describing, testing and authorizing systems or activities prior to or after a system is in operation. The C&A process is used extensively across the world.

Data integrity refers to maintaining and assuring the accuracy and consistency of [data](#) over its entire [life-cycle](#),^[1] and is a critical aspect to the design, implementation and usage of any system which stores, processes, or retrieves data. The term **data integrity** is broad in scope and may have widely different meanings depending on the specific context – even under the same general umbrella of [computing](#). This article provides only a broad overview of some of the different types and concerns of data integrity.

The **Brewer and Nash model** was constructed to provide [information security access controls](#) that can change dynamically. This security model, also known as the [Chinese wall](#) model, was designed to provide controls that mitigate [conflict of interest](#) in commercial organizations, and is built upon an [information flow model](#).

In the Brewer and Nash Model no information can flow between the subjects and objects in a way that would create a conflict of interest.

The **Bell–LaPadula Model** (abbreviated **BLP**) is a [state machine](#) model used for enforcing [access control](#) in government and military applications.^[1] It was developed by David Elliott Bell^[2] and Leonard J. LaPadula, subsequent to strong guidance from [Roger R. Schell](#) to formalize the [U.S. Department of Defense](#) (DoD) [multilevel security](#) (MLS) policy.^{[3][4][5]} The model is a formal [state transition model](#) of [computer security policy](#) that describes a set of access control rules which use security labels on objects and clearances for subjects. Security labels range from the most sensitive (e.g. "Top Secret"), down to the least sensitive (e.g., "Unclassified" or "Public").

The Bell–LaPadula model focuses on data [confidentiality](#) and controlled access to [classified information](#), in contrast to the [Biba Integrity Model](#) which describes rules for the protection of [data integrity](#). In this formal model, the entities in an [information system](#) are divided into subjects and objects. The notion of a "[secure state](#)" is defined, and it is proven that each state transition preserves security by moving from secure state to secure state, thereby [inductively](#) proving that the system satisfies the security objectives of the model. The Bell–LaPadula model is built on the concept of a [state machine](#) with a set of allowable states in a computer system. The transition from one state to another state is defined by transition functions.

A **digital signature** is a mathematical scheme for demonstrating the authenticity of a digital message or documents. A valid digital signature gives a recipient reason to believe that the message was created by a known sender, that the sender cannot deny having sent the message ([authentication](#) and [non-repudiation](#)), and that the message was not altered in transit ([integrity](#)). Digital signatures are commonly used for software distribution, financial transactions, and in other cases where it is important to detect forgery or tampering.

ISO 27001:2013 is an information security standard that was published on the 25th September 2013.^[1] It supersedes [ISO/IEC 27001:2005](#), and is published by the [International Organization for](#)

Standardization (ISO) and the **International Electrotechnical Commission** (IEC) under the joint ISO and IEC subcommittee, [ISO/IEC JTC 1/SC 27](#).^[2] It is a specification for an information security management system (ISMS). Organisations which meet the standard may gain an official certification issued by an independent and accredited certification body on successful completion of a formal audit process.

ISO/IEC 27002 is an [information security standard](#) published by the [International Organization for Standardization](#) (ISO) and by the [International Electrotechnical Commission](#) (IEC), titled *Information technology – Security techniques – Code of practice for information security management*.

ISO/IEC **27002**:2005 was developed from BS7799, published in the mid-1990s. The British Standard was adopted by ISO/IEC as ISO/IEC 17799:2000, revised in 2005, and renumbered (but otherwise unchanged) in 2007 to align with the other [ISO/IEC 27000-series](#) standards.^[clarification needed]

ISO/IEC 27002 provides [best practice](#) recommendations on information security management for use by those responsible for initiating, implementing or maintaining [information security management systems](#) (ISMS). Information security is defined within the standard in the context of the [C-I-A triad](#)

The **Common Criteria for Information Technology Security Evaluation** (abbreviated as **Common Criteria** or **CC**) is an [international standard](#) ([ISO/IEC 15408](#)) for [computer security](#) certification. It is currently in version 3.1 revision 4.^[1]

Common Criteria is a framework in which computer system users can *specify* their security *functional* and *assurance* requirements (SFRs and SARs respectively) through the use of Protection Profiles (PPs), vendors can then *implement* and/or make claims about the security attributes of their products, and testing laboratories can *evaluate* the products to determine if they actually meet the claims. In other words, Common Criteria provides assurance that the process of specification, implementation and evaluation of a computer security product has been conducted in a rigorous and standard and repeatable manner at a level that is commensurate with the target environment for use.^[2]

Common Criteria is used as the basis for a Government driven certification scheme and typically evaluations are conducted for the use of Federal Government agencies and critical infrastructure.

Internet Protocol Security (IPsec) is a [protocol suite](#) for secure [Internet Protocol](#) (IP) communications by [authenticating](#) and [encrypting](#) each [IP packet](#) of a communication session. IPsec includes protocols for establishing [mutual authentication](#) between agents at the beginning of the session and negotiation of [cryptographic keys](#) to be used during the session. IPsec can be used in protecting data flows between a pair of hosts (*host-to-host*), between a pair of security gateways (*network-to-network*), or between a security gateway and a host (*network-to-host*).^[1]

Internet Protocol security (IPsec) uses cryptographic security services to protect communications over Internet Protocol (IP) networks. IPsec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection.

IPsec is an end-to-end security scheme operating in the [Internet Layer](#) of the [Internet Protocol Suite](#), while some other Internet security systems in widespread use, such as [Transport Layer Security](#) (TLS) and [Secure Shell](#) (SSH), operate in the [upper layers](#) at the Application layer. Hence, only IPsec protects all application traffic over an IP network. Applications can be automatically secured by IPsec at the IP layer.

Encapsulating Security Payload^[edit]

Encapsulating Security Payload (ESP) is a member of the IPsec protocol suite. In IPsec it provides origin [authenticity](#), [integrity](#) and [confidentiality](#) protection of [packets](#). ESP also supports [encryption](#)-only and [authentication](#)-only configurations, but using encryption without authentication is strongly discouraged because it is insecure.^{[16][17][18]} Unlike [Authentication Header \(AH\)](#), ESP in transport mode does not provide integrity and authentication for the entire [IP packet](#). However, in [Tunnel Mode](#), where the entire original IP packet is [encapsulated](#) with a new packet header added, ESP protection is afforded to the whole inner IP packet (including the inner header) while the outer header (including any outer IPv4 options or IPv6 extension headers) remains unprotected. ESP operates directly on top of IP, using IP protocol number 50.^[15]

A **virtual private network (VPN)** extends a [private network](#) across a [public](#) network, such as the [Internet](#). It enables users to send and receive data across shared or public networks as if their computing devices were directly connected to the private network, and thus are benefiting from the functionality, security and management policies of the private network.^[1] A VPN is created by establishing a virtual [point-to-point](#) connection through the use of dedicated connections, virtual [tunneling protocols](#), or traffic [encryption](#).

A VPN spanning the Internet is similar to a [wide area network](#) (WAN). From a user perspective, the extended network resources are accessed in the same way as resources available within the private network.^[2] Traditional VPNs are characterized by a point-to-point topology, and they do not tend to support or connect [broadcast domains](#). Therefore, communication, software, and networking, which are based on [OSI layer](#) 2 and broadcast [packets](#), such as [NetBIOS](#) used in [Windows networking](#), may not be fully supported or work exactly as they would on a [local area network](#) (LAN). VPN variants, such as [Virtual Private LAN Service](#) (VPLS), and layer 2 tunneling protocols, are designed to overcome this limitation.

VPNs allow employees to securely access the corporate intranet while traveling outside the office. Similarly, VPNs securely connect geographically separated offices of an organization, creating one cohesive network. VPN technology is also used by individual Internet users to secure their [wireless](#) transactions, to circumvent geo-restrictions and censorship, and to connect to [proxy servers](#) for the purpose of protecting personal identity and location.

Database encryption is the process of converting [data](#), within a [database](#), in [plain textformat](#) into a meaningless [cipher text](#) by means of a suitable [algorithm](#). The database[encryption](#) protects the stored data. The act of encrypting a database also reduces the incentive for individuals to hack the aforementioned database as "meaningless" encrypted data is of little to no use for hackers.^[1]

There are multiple techniques and [technologies](#) available for database encryption. There are multiple styles of database presentation as well, and some of the [OEMs](#) provide[transparent data encryption](#). With TDE, some level of encryption can be achieved, but the[protocol](#) can cause degradation of performance in a database if it is not configured accurately.

Database encryption is done to encrypt sensitive data like [credit card numbers](#), [medical records](#), etc. on the [tables](#), [columns](#), or [rows](#) of a database. Database encryption requirements are sometimes governed by [regulation](#) or [business](#) or [data privacy laws](#)specific to different [countries](#).

Using database encryption can limit a [database administrator](#) in [copying](#) or seeing business-critical information.

Mutual Authentication is a security feature in which a client process must prove its identity to a server, and the server must prove its identity to the client, before any application traffic is sent over the client-to-server connection.

In computer networking, **ARP spoofing**, **ARP cache poisoning**, or**ARP poison routing**, is a technique by which an attacker sends ([spoofed](#))[Address Resolution Protocol](#) (ARP) messages onto a [local area network](#). Generally, the aim is to associate the attacker's [MAC address](#) with the [IP address](#) of another [host](#), such as the[default gateway](#), causing any traffic meant for that IP address to be sent to the attacker instead.

ARP spoofing may allow an attacker to intercept [data frames](#) on a network, modify the traffic, or stop all traffic. Often the attack is used as an opening for other attacks, such as [denial of service](#), [man in the middle](#), or [session hijacking](#) attacks.^[1]

The attack can only be used on networks that use the [Address Resolution Protocol](#), and is limited to local network segments.^[2]

DNS spoofing (or **DNS cache poisoning**) is a [computer hacking](#) attack, whereby data is introduced into a [Domain Name System](#) (DNS) [resolver's cache](#), causing the name server to return an incorrect [IP address](#), [diverting traffic](#) to the attacker's computer (or any other computer).

Message Validation

Message validation is used to protect your service from malformed messages and message parameters. Message schemas can be used to validate incoming messages, and custom validators can be used to validate parameter data before your service consumes it. Do not trust input from any source that the client can influence, such as cookies, headers, IP address, or the content of messages sent to your service. Do trust input from a database, the file system, or anything else outside the trust boundary of your service. Use message schemas and data validators to check for format, range, length, and type. Do not rely on client-side validation; make all security decisions based on server-side validation.

Cross-site scripting (XSS) is a type of [computer security vulnerability](#) typically found in [web applications](#). XSS enables attackers to [inject client-side script](#) into [web pages](#) viewed by other users. A cross-site scripting vulnerability may be used by attackers to [bypass access controls](#) such as the [same-origin policy](#). Cross-site scripting carried out on websites accounted for roughly 84% of all security vulnerabilities documented by [Symantec](#) as of 2007.^[1] Their effect may range from a petty nuisance to a significant security risk, depending on the sensitivity of the data handled by the vulnerable site and the nature of any security mitigation implemented by the site's owner.

In [computer security](#) and [programming](#), a **buffer overflow**, or **buffer overrun**, is an [anomaly](#) where a [program](#), while writing [data](#) to a [buffer](#), overruns the buffer's boundary and overwrites adjacent [memory](#) locations. This is a special case of the violation of [memory safety](#).

Buffer overflows can be triggered by inputs that are designed to [execute](#) code, or alter the way the program operates. This may result in erratic program behavior, including memory access errors, incorrect results, a [crash](#), or a breach of system security. Thus, they are the basis of many [software vulnerabilities](#) and can be maliciously [exploited](#).

[Programming languages](#) commonly associated with buffer overflows include [C](#) and [C++](#), which provide no built-in protection against accessing or overwriting data in any part of memory and do not automatically check that data written to an [array](#) (the built-in buffer type) is within the boundaries of that array. [Bounds checking](#) can prevent buffer overflows.

A buffer overflow occurs when [data](#) written to a buffer also corrupts data values in [memory addresses](#) adjacent to the destination buffer due to insufficient [bounds checking](#). This can occur

when copying data from one buffer to another without first checking that the data fits within the destination buffer.

A **rainbow table** is a [precomputed table](#) for reversing [cryptographic hash functions](#), usually for [cracking password](#) hashes. Tables are usually used in recovering a [plaintextpassword](#) up to a certain length consisting of a limited set of characters. It is a practical example of a [space/time trade-off](#), using less computer processing time and more storage than a [brute-force attack](#) which calculates a hash on every attempt, but more processing time and less storage than a simple [lookup table](#) with one entry per hash. Use of a [key derivation function](#) that employs a [salt](#) makes this attack infeasible.

All tokens contain some secret information that are used to prove identity. There are four different ways in which this information can be used:

Asynchronous password token for online banking.

1. Static password token. The device contains a password which is physically hidden (not visible to the possessor), but which is transmitted for each authentication. This type is vulnerable to [replay attacks](#).
2. Synchronous dynamic password token. A timer is used to rotate through various combinations produced by a cryptographic algorithm. The token and the authentication server must have synchronized clocks.
3. Asynchronous password token. A [one-time password](#) is generated without the use of a clock, either from a [one-time pad](#) or cryptographic algorithm.
4. [Challenge response](#) token. Using [public key cryptography](#), it is possible to prove possession of a private key without revealing that key. The authentication server encrypts a challenge (typically a random number, or at least data with some random parts) with a public key; the device proves it possesses a copy of the matching private key by providing the decrypted challenge.

Time-synchronized one-time passwords^{[[edit](#)]}

Time-synchronized one-time passwords change constantly at a set time interval, e.g. once per minute. To do this some sort of synchronization must exist between the [client](#)'s token and the authentication [server](#). For disconnected tokens this time-synchronization is done before the token is distributed to the [client](#). Other token types do the synchronization when the token is inserted into an [input device](#). The main problem with time-synchronized tokens is that they can, over time, become unsynchronized.^{[[citation needed](#)]} However, some such systems, such as RSA's [SecurID](#), allow the user to resynchronize the server with the token, sometimes by entering several consecutive

passcodes. Most also cannot have replaceable batteries and only last up to 5 years before having to be replaced - so there is additional cost.

Mathematical-algorithm-based one-time passwords^[edit]

Another type of one-time password uses a complex mathematical algorithm, such as a [hash chain](#), to generate a series of one-time passwords from a secret shared key. Each password is unguessable, even when previous passwords are known. The open source [OATH](#) algorithm is standardized; other algorithms are covered by U.S. [patents](#). Each password is observably unpredictable and independent on previous ones. Therefore, an adversary would be unable to guess what the next password may be, even with the knowledge of all previous passwords.

Non-repudiation refers to a state of affairs where the author of a statement will not be able to successfully challenge the authorship of the statement or validity of an associated [contract](#). The term is often seen in a legal setting wherein the [authenticity](#) of a signature is being challenged. In such an instance, the authenticity is being "repudiated". In a general sense *non-repudiation* involves associating actions or changes to a unique individual. For a secure area, for example, it may be desirable to implement a [key card access](#) system. Non-repudiation would be violated if it were not also a strictly enforced policy to prohibit sharing of the key cards and to immediately report lost or stolen cards. Otherwise determining who performed the action of opening the door cannot be trivially determined. Similarly, for computer accounts, the individual owner of the account must not allow others to use that account, especially, for instance, by giving away their account's password, and a policy should be implemented to enforce this. This prevents the owner of the account from denying actions performed by the account.^[1]

In digital security^[edit]

Regarding [digital security](#), the cryptological meaning and application of non-repudiation shifts to mean:^[2]

- A service that provides proof of the [integrity](#) and [origin of data](#).
- An authentication that can be asserted to be genuine with high assurance.

Proof of data integrity is typically the easiest of these requirements to accomplish. A data [hash](#), such as [SHA2](#), is usually sufficient to establish that the likelihood of data being undetectably changed is extremely low. Even with this safeguard, it is still possible to tamper with data in transit, either through a [man-in-the-middle attack](#) or [phishing](#). Due to this flaw, data integrity is best asserted when the recipient already possesses the necessary verification information.

The most common method of asserting the digital origin of data is through [digital certificates](#), a form of [public key infrastructure](#), to which [digital signatures](#) belong. Note that the public key scheme is not

used for [encryption](#) in this form, confidentiality is not achieved by signing a message with a private key (since anyone can obtain the public key to reverse the signature). Verifying the digital origin means that the certified/signed data can be, with reasonable certainty, trusted to be from somebody who possesses the private key corresponding to the signing certificate. If the key is not properly safeguarded by the original owner, digital forgery can become a major concern.

A **cryptographic hash function** is a [hash function](#) which is considered [practically impossible to invert](#), that is, to recreate the input data from its hash value alone. These one-way hash functions have been called "the workhorses of modern cryptography".^[1] The input data is often called the *message*, and the hash value is often called the *message digest* or simply the *digest*.

The ideal cryptographic hash function has four main properties:

- it is easy to compute the hash value for any given message
- it is [infeasible](#) to generate a message from its hash
- it is infeasible to modify a message without changing the hash
- it is infeasible to find two different messages with the same hash.

Cryptographic hash functions have many [information security](#) applications, notably in [digital signatures](#), [message authentication codes](#) (MACs), and other forms of [authentication](#). They can also be used as ordinary [hash functions](#), to index data in [hash tables](#), for [fingerprinting](#), to detect duplicate data or uniquely identify files, and as [checksums](#) to detect accidental data corruption. Indeed, in information security contexts, cryptographic hash values are sometimes called (*digital*) *fingerprints*, *checksums*, or just *hash values*, even though all these terms stand for more general functions with rather different properties and purposes.

In [computing](#), the **Challenge-Handshake Authentication Protocol (CHAP)** [authenticates](#) a user or network host to an authenticating entity. That entity may be, for example, an [Internet service provider](#). CHAP is specified in [RFC 1994](#).

CHAP provides protection against [replay attacks](#) by the peer through the use of an incrementally changing identifier and of a variable challenge-value. CHAP requires that both the client and server know the plaintext of the secret, although it is never sent over the network. Thus, CHAP provides better security as compared to [Password Authentication Protocol](#) (PAP) which is vulnerable for both these reasons. The [MS-CHAP](#) variant does not require either peer to know the plaintext and does not transmit it, but has been broken.^[1]

The **Point-to-Point Tunneling Protocol (PPTP)** is a method for implementing [virtual private networks](#). PPTP uses a control channel over [TCP](#) and a [GRE](#) tunnel operating to encapsulate [PPP](#) packets.

The PPTP specification does not describe [encryption](#) or [authentication](#) features and relies on the Point-to-Point Protocol being tunneled to implement security functionality. However, the most common PPTP implementation shipping with the [Microsoft Windows](#) product families implements various levels of authentication and encryption natively as standard features of the Windows PPTP stack. The intended use of this protocol is to provide security levels and remote access levels comparable with typical [VPN](#) products.

Simple Mail Transfer Protocol (SMTP) is an [Internet standard](#) for [electronic mail](#) (email) transmission. First defined by [RFC 821](#) in 1982, it was last updated in 2008 with the [Extended SMTP](#) additions by [RFC 5321](#)—which is the protocol in widespread use today.

SMTP by default uses [TCP port 25](#). The protocol for mail submission is the same, but uses port 587. SMTP connections secured by [SSL](#), known as [SMTPS](#), default to port 465 (nonstandard, but sometimes used for legacy reasons).

Although electronic [mail servers and other mail transfer agents](#) use SMTP to send and receive mail messages, user-level client mail applications typically use SMTP only for sending messages to a mail server for [relaying](#). For receiving messages, client applications usually use either [POP3](#) or [IMAP](#).

The **Hypertext Transfer Protocol (HTTP)** is an [application protocol](#) for distributed, collaborative, [hypermedia](#) information systems.^[1] HTTP is the foundation of data communication for the [World Wide Web](#).

[Hypertext](#) is structured text that uses logical links ([hyperlinks](#)) between [nodes](#) containing text. HTTP is the protocol to exchange or transfer hypertext.

The standards development of HTTP was coordinated by the [Internet Engineering Task Force](#) (IETF) and the [World Wide Web Consortium](#) (W3C), culminating in the publication of a series of [Requests for Comments](#) (RFCs). The first definition of HTTP/1.1, the version of HTTP in common use, occurred in [RFC 2068](#) in 1997, although this was obsoleted by [RFC 2616](#) i

IEEE 802.1X is an [IEEE Standard](#) for port-based [Network Access Control](#) (PNAC). It is part of the [IEEE 802.1](#) group of networking protocols. It provides an [authentication](#) mechanism to devices wishing to attach to a [LAN](#) or [WLAN](#).

IEEE 802.1X defines the encapsulation of the [Extensible Authentication Protocol](#) (EAP) over [IEEE 802](#),^{[1][2]} which is known as "EAP over LAN" or EAPOL.^[3] EAPOL was originally designed for [IEEE 802.3](#) Ethernet in 802.1X-2001, but was clarified to suit other IEEE 802 LAN technologies such as [IEEE 802.11](#) wireless and [Fiber Distributed Data Interface](#) (ISO 9314-2) in 802.1X-2004.^[4] The EAPOL protocol was also modified for use with [IEEE 802.1AE](#) ("MACsec") and [IEEE 802.1AR](#) (Secure Device Identity, DevID) in 802.1X-2010^{[5][6]} to support service identification and optional point to point encryption over the local LAN segment.

A **hypervisor** or **virtual machine monitor (VMM)** is a piece of computer software, firmware or hardware that creates and runs [virtual machines](#).

A computer on which a hypervisor is running one or more virtual machines is defined as a *host machine*. Each virtual machine is called a *guest machine*. The hypervisor presents the guest operating systems with a [virtual operating platform](#) and manages the execution of the guest operating systems. Multiple instances of a variety of operating systems may share the virtualized hardware resources.

HTTPS (also called **HTTP over TLS**,^{[1][2]} **HTTP over SSL**,^[3] and **HTTP Secure**^{[4][5]}) is a [protocol](#) for [secure](#) communication over a [computer network](#) which is widely used on the [Internet](#). HTTPS consists of communication over [Hypertext Transfer Protocol](#) (HTTP) within a connection encrypted by [Transport Layer Security](#) or its predecessor, [Secure Sockets Layer](#). The main motivation for HTTPS is [authentication](#) of the visited [website](#) and to protect the [privacy](#) and [integrity](#) of the exchanged data.

In its popular deployment on the internet, HTTPS provides authentication of the website and associated [web server](#) with which one is communicating, which protects against man-in-the-middle attacks. Additionally, it provides bidirectional [encryption](#) of communications between a client and server, which protects against [eavesdropping](#) and [tampering](#) with and/or forging the contents of the communication.^[6] In practice, this provides a reasonable guarantee that one is communicating with precisely the website that one intended to communicate with (as opposed to an impostor), as well as ensuring that the contents of communications between the user and site cannot be read or forged by any third party.

Historically, HTTPS connections were primarily used for payment transactions on the [World Wide Web](#), e-mail and for sensitive transactions in corporate information systems. In the late 2000s and early 2010s, HTTPS began to see widespread use for protecting page authenticity on all types of websites, securing accounts and keeping user communications, identity and [web browsing](#) private.

Pretty Good Privacy (PGP) is a [data encryption](#) and decryption [computer program](#) that provides [cryptographic privacy](#) and [authentication](#) for data communication. PGP is often used for signing, encrypting, and [decrypting](#) texts, e-mails, files, directories, and whole disk partitions and to increase the security of e-mail communications. It was created by [Phil Zimmermann](#) in 1991.^[1]

PGP and similar software follow the [OpenPGP](#) standard ([RFC 4880](#)) for encrypting and decrypting [data](#).

Cloud computing model for enabling ubiquitous, clear convenient, on-demand access to a shared pool of configurable computing resources.^{[1] [2]} Cloud computing and storage solutions provide users and enterprises with various capabilities to store and process their data in third-party [data centers](#).^[3] It relies on sharing of resources to achieve coherence and [economies of scale](#), similar to a utility (like the [electricity grid](#)) over a network.^[4] At the foundation of cloud computing is the broader concept of [converged infrastructure](#) and [shared services](#).

Cloud computing, or in simpler shorthand just "the cloud", also focuses on maximizing the effectiveness of the shared resources. Cloud resources are usually not only shared by multiple users but are also dynamically reallocated per demand. This can work for allocating resources to users. For example, a cloud computer facility that serves European users during European business hours with a specific application (e.g., email) may reallocate the same resources to serve North American users during North America's business hours with a different application (e.g., a web server). This approach helps maximize the use of computing power while reducing the overall cost of resources by using less power, air conditioning, rack space, etc. to maintain the system. With cloud computing, multiple users can access a single server to retrieve and update their data without purchasing licenses for different applications.

Infrastructure as a service (IaaS)^{[[edit](#)]}

See also: [Category:Cloud infrastructure](#)

In the most basic cloud-service model - and according to the IETF (Internet Engineering Task Force) - providers of IaaS offer computers – physical or (more often) virtual machines – and other resources. IaaS refers to online services that abstract user from the detail of infrastructure like

physical computing resources, location, data partitioning, scaling, security, backup etc. A [hypervisor](#), such as [Xen](#), [Oracle VirtualBox](#), [KVM](#), [VMware ESX/ESXi](#), or [Hyper-V](#) runs the virtual machines as guests. Pools of hypervisors within the cloud operational system can support large numbers of virtual machines and the ability to scale services up and down according to customers' varying requirements. IaaS clouds often offer additional resources such as a virtual-machine [disk-image](#) library, raw [block storage](#), file or [object storage](#), firewalls, load balancers, IP addresses, [virtual local area networks](#) (VLANs), and software bundles.^[56] IaaS-cloud providers supply these resources on-demand from their large pools of equipment installed in [data centers](#). For [wide-area](#) connectivity, customers can use either the Internet or [carrier clouds](#) (dedicated [virtual private networks](#)).

To deploy their applications, cloud users install operating-system images and their application software on the cloud infrastructure. In this model, the cloud user patches and maintains the operating systems and the application software. Cloud providers typically bill IaaS services on a utility computing basis: cost reflects the amount of resources allocated and consumed.^{[57][58][59][60]}

Platform as a service (PaaS)^[edit]

Main article: [Platform as a service](#)

See also: [Category:Cloud platforms](#)

PaaS vendors offer a development environment to application developers. The provider typically develops toolkit and standards for development and channels for distribution and payment. In the PaaS models, cloud providers deliver a [computing platform](#), typically including operating system, programming-language execution environment, database, and web server. Application developers can develop and run their software solutions on a cloud platform without the cost and complexity of buying and managing the underlying hardware and software layers. With some PaaS offers like [Microsoft Azure](#) and [Google App Engine](#), the underlying computer and storage resources scale automatically to match application demand so that the cloud user does not have to allocate resources manually. The latter^[which?] has also been proposed by an architecture aiming to facilitate real-time in cloud environments.^[61]^[*need quotation to verify*] Even more specific application types can be provided via PaaS, such as media encoding as provided by services like bitcodin.com^[62] or media.io.^[63]

Some integration and data management providers have also embraced specialized applications of PaaS as delivery models for data solutions. Examples include **iPaaS** and **dPaaS**. iPaaS (Integration Platform as a Service) enables customers to develop, execute and govern integration flows.^[64] Under the iPaaS integration model, customers drive the development and deployment of integrations without installing or managing any hardware or middleware.^[65] dPaaS (Data Platform as a Service) delivers integration—and data-management—products as a fully managed service.^[66] Under the dPaaS model, the PaaS provider, not the customer, manages the development and execution of

data solutions by building tailored data applications for the customer. dPaaS users retain transparency and control over data through [data-visualization](#) tools.^[67]

Software as a service (SaaS)^[edit]

Main article: [Software as a service](#)

In the software as a service (SaaS) model, users gain access to application software and databases. Cloud providers manage the infrastructure and platforms that run the applications. SaaS is sometimes referred to as "on-demand software" and is usually priced on a pay-per-use basis or using a subscription fee.^[*citation needed*]

In the SaaS model, cloud providers install and operate application software in the cloud and cloud users access the software from cloud clients. Cloud users do not manage the cloud infrastructure and platform where the application runs. This eliminates the need to install and run the application on the cloud user's own computers, which simplifies maintenance and support. Cloud applications differ from other applications in their scalability—which can be achieved by cloning tasks onto multiple [virtual machines](#) at run-time to meet changing work demand.^[68] [Load balancers](#) distribute the work over the set of virtual machines. This process is transparent to the cloud user, who sees only a single access-point. To accommodate a large number of cloud users, cloud applications can be [multitenant](#), meaning that any machine may serve more than one cloud-user organization.

The pricing model for SaaS applications is typically a monthly or yearly flat fee per user,^[69] so prices become scalable and adjustable if users are added or removed at any point.^[70]

Proponents claim that SaaS gives a [business](#) the potential to reduce IT operational costs by outsourcing hardware and software maintenance and support to the cloud provider. This enables the business to reallocate IT operations costs away from hardware/software spending and from personnel expenses, towards meeting other goals. In addition, with applications hosted centrally, updates can be released without the need for users to install new software. One drawback of SaaS comes with storing the users' data on the cloud provider's server. As a result,^[*citation needed*] there could be unauthorized access to the data. For this reason, users are increasingly^[*quantify*] adopting intelligent third-party [key-management](#) systems to help secure their data.

A **passive infrared sensor (PIR sensor)** is an electronicsensor that measures [infrared](#) (IR) light radiating from objects in its field of view. They are most often used in [PIR-based motion detectors](#). All objects with a temperature above [absolute zero](#) emit [heat](#) energy in the form of radiation. Usually this radiation is invisible to the [human eye](#) because it radiates at infrared wavelengths, but it can be detected by electronic devices designed for such a purpose.

The term *passive* in this instance refers to the fact that PIR devices do not generate or radiate any energy for detection purposes. They work entirely by detecting the energy given off by other

objects.^[1] PIR sensors don't detect or measure "heat"; instead they detect the infrared radiation emitted or reflected from an object.

